Practitioner's Docket No. FORE-56 **PATENT**

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Box Patent Application
Assistant Commissioner for Patents
Washington, D.C. 20231

NEW APPLICATION TRANSMITTAL

Transmitted herewith for filing is the patent application of

Inventor(s): Seth Redmore

WARNING: *37 C.F.R. § 1.41(a)(1) points out:*

"(a) A patent is applied for in the name or names of the actual inventor or inventors.

"(1) The inventorship of a nonprovisional application is that inventorship set forth in the oath or declaration as prescribed by § 1.63, except as provided for in § 1.53(d)(4) and § 1.63(d). If an oath or declaration as prescribed by § 1.63 is not filed during the pendency of a nonprovisional application, the inventorship is that inventorship set forth in the application papers filed pursuant to § 1.53(b), unless a petition under this paragraph accompanied by the fee set forth in § 1.17(i) is filed supplying or changing the name or names of the inventor or inventors."

For (title): HARDWARE BASED SECURITY GROUPS, FIREWALL LOAD SHARING, AND FIREWALL REDUNDANCY

CERTIFICATION UNDER 37 C.F.R. 1.10*
(Express Mail label number is mandatory.)
(Express Mail certification is optional.)

I hereby certify that this New Application Transmittal and the documents referred to as attached therein are being deposited with the United States Postal Service on this date July 16, 1999, in an envelope as "Express Mail Post Office to Addressee," mailing Label Number EL262551202US, addressed to the: Assistant Commissioner for Patents, Washington, D.C. 20231.

Tracey L. Milka
(type or print name of person mailing paper)

Tracey L. Milka
Signature of person mailing paper

WARNING: *Certificate of mailing (first class) or facsimile transmission procedures of 37 C.F.R. 1.8 cannot be used to obtain a date of mailing or transmission for this correspondence.*

****WARNING:*** *Each paper or fee filed by "Express Mail"* ***must*** *have the number of the "Express Mail" mailing label placed thereon prior to mailing. 37 C.F.R. 1.10(b).*
"Since the filing of correspondence under § 1.10 without the Express Mail mailing label thereon is an oversight that can be avoided by the exercise of reasonable care, requests for waiver of this requirement will ***not*** *be granted on petition." Notice of Oct. 24, 1996, 60 Fed. Reg. 56,439, at 56,442.*

1. Type of Application

This new application is for a(n)

(check one applicable item below)

☒ Original (nonprovisional)

☐ Design

☐ Plant

WARNING: ***Do not*** *use this transmittal for a completion in the U.S. of an International Application under 35 U.S.C. 371(c)(4), unless the International Application is being filed as a divisional, continuation or continuation-in-part application.*

WARNING: ***Do not*** *use this transmittal for the filing of a provisional application.*

NOTE: If one of the following 3 items apply, then complete and attach ADDED PAGES FOR NEW APPLICATION TRANSMITTAL WHERE BENEFIT OF A PRIOR U.S. APPLICATION CLAIMED and a NOTIFICATION IN PARENT APPLICATION OF THE FILING OF THIS CONTINUATION APPLICATION.

☐ Divisional.

☐ Continuation.

☐ Continuation-in-part (C-I-P).

2. Benefit of Prior U.S. Application(s) (35 U.S.C. 119(e), 120, or 121)

NOTE: A nonprovisional application may claim an invention disclosed in one or more prior filed copending nonprovisional applications or copending international applications designating the United States of America. In order for a nonprovisional application to claim the benefit of a prior filed copending nonprovisional application or copending international application designating the United States of America, each prior application must name as an inventor at least one inventor named in the later filed nonprovisional application and disclose the named inventor's invention claimed in at least one claim of the later filed nonprovisional application in the manner provided by the first paragraph of 35 U.S.C. 112. Each prior application must also be:

(i) An international application entitled to a filing date in accordance with PCT Article 11 and designating the United States of America; or

(ii) Complete as set forth in § 1.51(b); or

(iii) Entitled to a filing date as set forth in § 1.53(b) or § 1.53(d) and include the basic filing fee set forth in § 1.16; or

(iv) Entitled to a filing date as set forth in § 1.53(b) and have paid therein the processing and retention fee set forth in § 1.21(l) within the time period set forth in § 1.53(f).

37 C.F.R. § 1.78(a)(1).

NOTE: If the new application being transmitted is a divisional, continuation or a continuation-in-part of a parent case, or where the parent case is an International Application which designated the U.S., or benefit of a prior provisional application is claimed, then check the following item and complete and attach ADDED PAGES FOR NEW APPLICATION TRANSMITTAL WHERE BENEFIT OF PRIOR U.S. APPLICATION(S) CLAIMED.

WARNING: *If an application claims the benefit of the filing date of an earlier filed application under 35 U.S.C. 120, 121 or 365(c), the 20–year term of that application will be based upon the filing date of the earliest U.S. application that the application makes reference to under 35 U.S.C. 120, 121 or 365(c). (35 U.S.C. 154(a)(2) does not take into account, for the determination of the patent term, any application on which priority is claimed under 35 U.S.C. 119, 365(a) or 365(b).) For a c–i–p application, applicant should review whether any claim in the patent that will issue is supported by an earlier application and, if not, the applicant should consider canceling the reference to the earlier filed application. The term of a patent is not based on a claim-by-claim approach. See Notice of April 14, 1995, 60 Fed. Reg. 20,195, at 20,205.*

WARNING: *When the last day of pendency of a provisional application falls on a Saturday, Sunday, or Federal holiday within the District of Columbia, any nonprovisional application claiming benefit of the provisional application* ***must*** *be filed prior to the Saturday, Sunday, or Federal holiday within the District of Columbia. See 37 C.F.R. § 1.78(a)(3).*

□ The new application being transmitted claims the benefit of prior U.S. application(s). Enclosed are ADDED PAGES FOR NEW APPLICATION TRANSMITTAL WHERE BENEFIT OF PRIOR U.S. APPLICATION(S) CLAIMED.

3. Papers Enclosed

A. Required for filing date under 37 C.F.R. § 1.53(b) (Regular) or 37 C.F.R. § 1.153 (Design) Application

23 Pages of specification

6 Pages of claims

2 Sheets of drawing

□ formal

☒ informal

B. Other Papers Enclosed

1 Pages of Abstract

0 Other

WARNING: ***DO NOT*** *submit original drawings. A high quality copy of the drawings should be supplied when filing a patent application. The drawings that are submitted to the Office must be on strong, white, smooth, and non-shiny paper and meet the standards according to § 1.84. If corrections to the drawings are necessary, they should be made to the original drawing and a high-quality copy of the corrected original drawing then submitted to the Office. Only one copy is required or desired. For comments on proposed then-new 37 CFR 1.84, see Notice of March 9, 1988 (1990 O.G. 57-62).*

NOTE: "Identifying indicia, if provided, should include the application number or the title of the invention, inventor's name, docket number (if any), and the name and telephone number of a person to call if the Office is unable to match the drawings to the proper application. This information should be placed on the back of each sheet of drawing a minimum distance of 1.5 cm. (5/8 inch) down from the top of the page . . ." 37 C.F.R. 1.84(c)).

(complete the following, if applicable)

□ The enclosed drawing(s) are photograph(s), and there is also attached a "PETITION TO ACCEPT PHOTOGRAPH(S) AS DRAWING(S)." 37 C.F.R. 1.84(b).

4. Additional papers enclosed

□ Preliminary Amendment

□ Information Disclosure Statement (37 C.F.R. 1.98)

□ Form PTO–1449 (PTO/SB/08A and 08B)

□ Citations

□ Declaration of Biological Deposit

□ Submission of "Sequence Listing," computer readable copy and/or amendment pertaining thereto for biotechnology invention containing nucleotide and/or amino acid sequence.

□ Authorization of Attorney(s) to Accept and Follow Instructions from Representative

□ Special Comments

□ Other

5. Declaration or oath

NOTE: *A newly executed declaration is not required in a continuation or divisional application provided that the prior nonprovisional application contained a declaration as required, the application being filed is by all or fewer than all the inventors named in the prior application, there is no new matter in the application being filed, and a copy of the executed declaration filed in the prior application (showing the signature or an indication thereon that it was signed) is submitted. The copy must be accompanied by a statement requesting deletion of the names of person(s) who are not inventors of the application being filed. If the declaration in the prior application was filed under § 1.47, then a copy of that declaration must be filed accompanied by a copy of the decision granting § 1.47 status or, if a nonsigning person under § 1.47 has subsequently joined in a prior application, then a copy of the subsequently executed declaration must be filed. See 37 C.F.R. §§ 1.63(d).*

□ Enclosed

Executed by

(check **all** *applicable boxes)*

□ inventor(s).

□ legal representative of inventor(s).
37 CFR 1.42 or 1.43.

□ joint inventor or person showing a proprietary
interest on behalf of inventor who refused to sign
or cannot be reached.

□ This is the petition required by 37 CFR 1.47 and the statement required by 37 CFR 1.47 is also attached. *See* item 13 below for fee.

☒ Not Enclosed.

NOTE: *Where the filing is a completion in the U.S. of an International Application or where the completion of the U.S. application contains subject matter in addition to the International Application, the application may be treated as a continuation or continuation-in-part, as the case may be, utilizing ADDED PAGE FOR NEW APPLICATION TRANSMITTAL WHERE BENEFIT OF PRIOR U.S. APPLICATION CLAIMED.*

□ Application is made by a person authorized under 37 C.F.R. 1.41(c) on behalf of *all* the above named inventor(s).

(The declaration or oath, along with the surcharge required by 37 CFR 1.16(e) can be filed subsequently).

NOTE: *It is important that all the correct inventor(s) are named for filing under 37 CFR 1.41(c) and 1.53(b).*

□ Showing that the filing is authorized.
(not required unless called into question. 37 CFR 1.41(d))

6. Inventorship Statement

WARNING: *If the named inventors are each not the inventors of all the claims an explanation, including the ownership of the various claims at the time the last claimed invention was made, should be submitted.*

The inventorship for all the claims in this application are:

☒ The same.

or

□ Not the same. An explanation, including the ownership of the various claims at the time the last claimed invention was made,

□ is submitted.

□ will be submitted.

7. Language

NOTE: *An application including a signed oath or declaration may be filed in a language other than English. An English translation of the non-English language application and the processing fee of $130.00 required by 37 CFR 1.17(k) is required to be filed with the application, or within such time as may be set by the Office. 37 CFR 1.52(d).*

☒ English

□ Non-English

□ The attached translation includes a statement that the translation is accurate. 37 C.F.R. 1.52(d).

8. Assignment

☒ An assignment of the invention to FORE Systems, Inc.

□ is attached. A separate □ "COVER SHEET FOR ASSIGNMENT (DOCUMENT) ACCOMPANYING NEW PATENT APPLICATION" or □ FORM PTO 1595 is also attached.

☒ will follow.

NOTE: *"If an assignment is submitted with a new application, send two separate letters-one for the application and one for the assignment." Notice of May 4, 1990 (1114 O.G. 77-78).*

WARNING: *A newly executed "CERTIFICATE UNDER 37 CFR 3.73(b)" must be filed when a continuation-in-part application is filed by an assignee. Notice of April 30, 1993, 1150 O.G. 62-64.*

9. Certified Copy

Certified copy(ies) of application(s)

Country	Appln. No.	Filed
Country	Appln. No.	Filed
Country	Appln. No.	Filed

from which priority is claimed

□ is (are) attached.

□ will follow.

NOTE: *The foreign application forming the basis for the claim for priority must be referred to in the oath or declaration. 37 CFR 1.55(a) and 1.63.*

NOTE: *This item is for any foreign priority for which the application being filed directly relates. If any parent U.S. application or International Application from which this application claims benefit under 35 U.S.C. 120 is itself entitled to priority from a prior foreign application, then complete item 18 on the ADDED PAGES FOR NEW APPLICATION TRANSMITTAL WHERE BENEFIT OF PRIOR U.S. APPLICATION(S) CLAIMED.*

10. Fee Calculation (37 C.F.R. 1.16)

A. ☒ Regular application

CLAIMS AS FILED

	Number filed	Number Extra		Rate	Basic Fee 37 C.F.R. 1.16(a) $790.00
Total Claims (37 CFR 1.16(c))	20– 20 =	0	×	$ 22.00	0.00
Independent Claims (37 CFR 1.16(b))	2 – 3 =	0	×	$ 82.00	0.00
Multiple dependent claim(s), if any (37 CFR 1.16(d))			+	$270.00	

□ Amendment cancelling extra claims is enclosed.

□ Amendment deleting multiple-dependencies is enclosed.

□ Fee for extra claims is not being paid at this time.

NOTE: *If the fees for extra claims are not paid on filing they must be paid or the claims cancelled by amendment, prior to the expiration of the time period set for response by the Patent and Trademark Office in any notice of fee deficiency. 37 CFR 1.16(d).*

Filing Fee Calculation $ 760.00

B. □ Design application
($330.00—37 CFR 1.16(f))

Filing Fee Calculation $__________

C. □ Plant application
($540.00—37 CFR 1.16(g))

Filing fee calculation $__________

11. Small Entity Statement(s)

□ Statement(s) that this is a filing by a small entity under 37 CFR 1.9 and 1.27 is (are) attached.

WARNING: *"Status as a small entity must be specifically established in each application or patent in which the status is available and desired. Status as a small entity in one application or patent does not affect any other application or patent, including applications or patents which are directly or indirectly dependent upon the application or patent in which the status has been established. The refiling of an application under § 1.53 as a continuation, division, or continuation-in-part (including a continued prosecution application under § 1.53(d)), or the filing of a reissue application requires a new determination as to continued entitlement to small entity status for the continuing or reissue application. A nonprovisional application claiming benefit under 35 U.S.C. 119(e), 120, 121, or 365(c) of a prior application, or a reissue application may rely on a statement filed in the prior application or in the patent if the nonprovisional application or the reissue application includes a reference to the statement in the prior application or in the patent or includes a copy of the statement in the prior application or in the patent and status as a small entity is still proper and desired. The payment of the small entity basic statutory filing fee will be treated as such a reference for purposes of this section." 37 C.F.R. § 1.28(a)(2).*

(complete the following, if applicable)

□ Status as a small entity was claimed in prior application
_______ / ________________, filed on ________________, from which benefit is being claimed for this application under:

35 U.S.C. □ 119(e),
□ 120,
□ 121,
□ 365(c),

and which status as a small entity is still proper and desired.

□ A copy of the statement in the prior application is included.

Filing Fee Calculation (50% of **A, B** or **C** above)

$________________________

NOTE: *Any excess of the full fee paid will be refunded if small entitiy status is established and a refund request are filed within 2 months of the date of timely payment of a full fee. The two-month period is not extendable under § 1.136. 37 CFR 1.28(a).*

12. Request for International-Type Search (37 C.F.R. 1.104(d))

(complete, if applicable)

□ Please prepare an international-type search report for this application at the time when national examination on the merits takes place.

13. Fee Payment Being Made at This Time

□ Not Enclosed

□ No filing fee is to be paid at this time.
(This and the surcharge required by 37 C.F.R. 1.16(e) can be paid subsequently.)

☒ Enclosed

☒ Filing fee	$ 760.00
□ Recording assignment ($40.00; 37 C.F.R. 1.21(h)) (See attached "COVER SHEET FOR ASSIGNMENT ACCOMPANYING NEW APPLICATION".)	$ ________
□ Petition fee for filing by other than all the inventors or person on behalf of the inventor where inventor refused to sign or cannot be reached ($130.00; 37 C.F.R. 1.47 and 1.17(i))	$ ________
□ For processing an application with a specification in a non-English language ($130.00; 37 C.F.R. 1.52(d) and 1.17(k))	$ ________
□ Processing and retention fee ($130.00; 37 C.F.R. 1.53(d) and 1.21(l))	$ ________
□ Fee for international-type search report ($40.00; 37 C.F.R. 1.21(e))	$ ________

NOTE: 37 CFR 1.21(l) establishes a fee for processing and retaining any application that is abandoned for failing to complete the application pursuant to 37 CFR 1.53(f) and this, as well as the changes to 37 CFR 1.53 and 1.78(a)(1), indicate that in order to obtain the benefit of a prior U.S. application, either the basic filing fee must be paid, or the processing and retention fee of § 1.21(l) must be paid, within 1 year from notification under § 53(f).

Total fees enclosed $ 760.00

14. Method of Payment of Fees

☒ Check in the amount of $ 760.00

☐ Charge Account No. ______________ in the amount of $__________.

A duplicate of this transmittal is attached.

NOTE: Fees should be itemized in such a manner that it is clear for which purpose the fees are paid. 37 CFR 1.22(b).

15. Authorization to Charge Additional Fees

WARNING: *If no fees are to be paid on filing, the following items should not be completed.*

WARNING: *Accurately count claims, especially multiple dependent claims, to avoid unexpected high charges, if extra claim charges are authorized.*

☒ The Commissioner is hereby authorized to charge the following additional fees by this paper and during the entire pendency of this application to Account No. 19-0737 :

☒ 37 C.F.R. 1.16(a), (f) or (g) (filing fees)

☒ 37 C.F.R. 1.16(b), (c) and (d) (presentation of extra claims)

NOTE: Because additional fees for excess or multiple dependent claims not paid on filing or on later presentation must only be paid or these claims cancelled by amendment prior to the expiration of the time period set for response by the PTO in any notice of fee deficiency (37 CFR 1.16(d)), it might be best not to authorize the PTO to charge additional claim fees, except possibly when dealing with amendments after final action.

☐ 37 C.F.R. 1.16(e) (surcharge for filing the basic filing fee and/or declaration on a date later than the filing date of the application)

☐ 37 C.F.R. §§ 1.17(a)(1)–(5) (extension fees pursuant to § 1.136(a)).

☐ 37 C.F.R. 1.17 (application processing fees)

NOTE: ". . .A written request may be submitted in an application that is an authorization to treat any concurrent or future reply, requiring a petition for an extension of time under this paragraph for its timely submission, as incorporating a petition for extension of time for the appropriate length of time. An authorization to charge all required fees, fees under § 1.17, or all required extension of time fees will be treated as a constructive petition for an extension of time in any concurrent or future reply requiring a petition for an extension of time under this paragraph for its timely submission. Submission of the fee set forth in § 1.17(a) will also be treated as a constructive petition for an extension of time in any concurrent reply requiring a petition for an extension of time under this paragraph for its timely submission." 37 C.F.R. § 1.136(a)(3).

☐ 37 C.F.R. 1.18 (issue fee at or before mailing of Notice of Allowance, pursuant to 37 C.F.R. 1.311(b))

NOTE: Where an authorization to charge the issue fee to a deposit account has been filed before the mailing of a Notice of Allowance, the issue fee will be automatically charged to the deposit account at the time of mailing the notice of allowance. 37 CFR 1.311(b).

NOTE: 37 CFR 1.28(b) requires "Notification of any change in status resulting in loss of entitlement to small entity status must be filed in the application . . . prior to paying, or at the time of paying, . . . issue fee." From the wording of 37 CFR 1.28(b), (a) notification of change of status must be made even if the fee is paid as "other than a small entity" and (b) no notification is required if the change is to another small entity.

16. Instructions as to Overpayment

NOTE: ". . . Amounts of twenty-five dollars or less will not be returned unless specifically requested within a reasonable time, nor will the payer be notified of such amounts; amounts over twenty-five dollars may be returned by check or, if requested, by credit to a deposit account." 37 C.F.R. § 1.26(a).

☒ Credit Account No. 19-0737

☐ Refund

Ansel Schwartz

SIGNATURE OF PRACTITIONER

Reg. No. 30,587

Ansel M. Schwartz

(type or print name of attorney)

Tel. No. (412) 621-9222

One Sterling Plaza

P.O. Address

201 N. Craig Street, Suite 304
Pittsburgh, PA 15213

Customer No.

□ **Incorporation by reference of added pages**

(check the following item if the application in this transmittal claims the benefit of prior U.S. application(s) (including an international application entering the U.S. stage as a continuation, divisional or C-I-P application) and complete and attach the ADDED PAGES FOR NEW APPLICATION TRANSMITTAL WHERE BENEFIT OF PRIOR U.S. APPLICATION(S) CLAIMED)

□ Plus Added Pages for New Application Transmittal Where Benefit of Prior U.S. Application(s) Claimed

Number of pages added ____________

□ Plus Added Pages for Papers Referred to in Item 4 Above

Number of pages added ____________

□ Plus added pages deleting names of inventor(s) named in prior application(s) who is/are no longer inventor(s) of the subject matter claimed in this application.

Number of pages added ____________

□ Plus "Assignment Cover Letter Accompanying New Application"

Number of pages added ____________

☒ **Statement Where No Further Pages Added**

(if no further pages form a part of this Transmittal, then end this Transmittal with this page and check the following item)

☒ This transmittal ends with this page.

HARDWARE BASED SECURITY GROUPS, FIREWALL LOAD SHARING, AND FIREWALL REDUNDANCY

FIELD OF THE INVENTION

The present invention is related to firewalls. More specifically, the present invention is related to firewalls connected by a switch to destinations.

BACKGROUND OF THE INVENTION

Certain types of network traffic in certain applications need to be inspected by a device which has lots of flexibility; e.g. a processor. The primary example of this is security, whereby every flow coming through a secured point (firewall) must be inspected to some greater or lesser degree. Hardware-based systems tend to be very fast, but don't deal well with very complex operations. Hence, software-based (processor-based) systems are still the norm, even with all of their concomitant performance problems. Fore instance, see U.S. Patent No. 5,699,513, incorporated by reference herein.

Even in a system where a single processor is fast enough, if that processor dies, then the whole system grinds to a standstill. This is highly undesirable in a mission critical application (or in any application for that matter).

SUMMARY OF THE INVENTION

The present invention pertains to a secure telecommunications system. The system comprises a network on which traffic travels. The system comprises a switch connected to the network. The system comprises a first inspection engine connected to the switch, which receives traffic from the switch, processes the traffic to determine whether it is desired traffic or undesired traffic, which prevents undesired traffic from passing through it and which sends desired traffic back to the switch. The system comprises a second inspection engine connected to the switch, which receives traffic from the switch, processes the traffic to determine whether it is desired traffic or undesired traffic, which prevents undesired traffic from passing through it and which sends desired traffic back to the switch. The system comprises a first destination connected to the switch which receives desired traffic from the switch that has been processed by the first inspection engine. The system comprises a second destination connected to the switch which receives desired traffic from the switch that has been processed by the second inspection engine.

The present invention pertains to a method for sending traffic over a secure telecommunications system. The method comprises the steps of receiving traffic from a network at a switch connected to the network. Then there is the step of directing traffic to a first inspection engine connected to the switch and to a second inspection engine connected to the switch. Next there is the step of receiving traffic at the first inspection engine. Then there is the step of processing traffic received at the first inspection engine to determine whether it is desired traffic or undesired traffic. Next there is the step of sending the desired traffic back to the switch from the first inspection engine and discarding undesired traffic from the first inspection engine. Then there is the step of transferring desired traffic received by the switch from the first inspection engine to a first destination. Next there is the step of processing traffic received at the second inspection engine to determine whether it is desired traffic or undesired traffic. Then there is the step of sending the desired traffic back to the switch from the second inspection engine and discarding undesired traffic from the second inspection engine. Next there is the step of transferring desired traffic received by the switch from the second inspection engine to a second destination.

BRIEF DESCRIPTION OF THE DRAWINGS

In the accompanying drawings, the preferred embodiment of the invention and preferred methods of practicing the invention are illustrated in which:

Figure 1 is a schematic representation of a system of the present invention.

Figure 2 is a schematic representation of another embodiment of a system of the present invention.

Figure 3 is a schematic representation of another embodiment of a system of the present invention.

DETAILED DESCRIPTION

Referring now to the drawings wherein like reference numerals refer to similar or identical parts throughout the several views, and more specifically to figure 1 thereof, there is shown a secure telecommunications system 10. The system 10 comprises a network 12 on which traffic travels. The system 10 comprises a switch 14 connected to the network 12. The system 10 comprises a first inspection engine 16 connected to the switch 14, which receives traffic from the switch 14, processes the traffic to determine whether it is desired traffic or undesired traffic, which prevents undesired traffic from passing through it and which sends desired traffic back to the switch 14. The system 10 comprises a second inspection engine 18 connected to the switch 14, which receives traffic from the switch 14, processes the traffic to determine whether it is desired traffic or undesired traffic, which prevents undesired traffic from passing through it and which sends desired traffic back to the switch 14. The system 10 comprises a first destination 20 connected to the switch 14 which receives desired traffic from the switch 14 that has been processed by the first inspection engine 16. The system 10 comprises a second destination 22 connected to the switch 14 which receives desired traffic from the switch 14 that has been processed by the second inspection engine 18.

Preferably, the first inspection engine 16 includes a first firewall processing engine 24 and the second inspection engine 18 includes a second firewall processing engine 26. The switch 14 preferably has a first port 28 and a second port 30 connected to the network 12 which receives traffic from the network 12. The switch 14 directing traffic received at the first port 28 to the first firewall processing engine 24 and directing traffic received at the second port 30 to the second firewall processing engine 26.

Preferably, the system 10 includes N additional firewall processing engines 32 connected to the switch 14 besides the first firewall processing engine 24 and the second firewall processing engine 26 so there are a total of N+2 firewall processing engines 32, where N is greater than or equal to 1 and is an integer. The switch 14 preferably has N additional ports 34 besides the first port 28 and the second port 30, wherein each port is connected to a corresponding firewall processing engine.

Preferably, the switch 14 is configured into security groups 36, as shown in figure 2, with at least one of the N+2 firewall processing engines 32 serving each security group 36. The switch 14 preferably load-shares traffic for each security group 36 across corresponding firewall processing engines 32 serving the corresponding security group 36. Preferably, the switch 14 rebalances traffic for a security group 36 when one of the firewall processing engines 32 serving the security group 36 fails across the other firewall processing engines 32 serving the security group 36. The switch 14 preferably is scalable to allow for adding firewall processing engines 32.

Preferably, the traffic includes bits and wherein the firewall processing engines 32 serving a first security group 36a of the security groups 36 encrypt greater than 1 Gbps of traffic. The network 12 preferably includes the Internet, and the first destination 20 includes a first web server and the second destination 22 includes a second web server. Preferably, the Internet includes a LAN.

The present invention pertains to a method for sending traffic over a secure telecommunications system 10. The method comprises the steps of receiving traffic from a network 12 at a switch 14 connected to the network 12. Then there is the step of directing traffic to a first inspection engine 16 connected to the switch 14 and to a second inspection engine 18 connected to the switch 14. Next there is the step of receiving traffic at the first inspection engine 16. Then there is the step of processing traffic received at the first inspection engine 16 to determine whether it is desired traffic or undesired traffic. Next there is the step of sending the desired traffic back to the switch 14 from the first inspection engine 16 and discarding undesired traffic from the first inspection engine 16. Then there is the step of transferring desired traffic received by the switch 14 from the first inspection engine 16 to a first destination 20. Next there is the step of processing traffic received at the second inspection engine 18 to determine whether it is desired traffic or undesired traffic. Then there is the step of sending the desired traffic back to the switch 14 from the second inspection engine 18 and discarding undesired traffic from the second inspection engine 18. Next there is the step of transferring desired traffic received by the switch 14 from the second inspection engine 18 to a second destination 22.

Preferably, the first and second inspection engines 18 include a first firewall processing engine 24 and a second firewall processing engine 26, respectively, and wherein the directing traffic step includes the step of directing traffic to the first firewall processing engine 24 and second firewall processing engine 26 and to a third firewall processing engine and a forth firewall processing engine. The switch 14 is preferably configured into a first security group 36a and a second security group 36b, and the receiving step includes the step of receiving traffic at the first security group 36a.

Preferably, the directing step includes the step of directing the traffic from the first security group 36a of the switch 14 to the first, third and fourth firewall processing engines 32 which serve the first security group 36a of the switch 14, and directing traffic to the second firewall processing engine 26 serving the second security group 36b of the switch 14. The receiving step preferably includes the step of receiving traffic from the first security group 36a at a first port 28 of the switch 14 and receiving traffic for the second security group 36b at a second port 30 of the switch 14. Preferably, the directing the traffic from the first security group 36a includes the step of load-sharing by the switch 14 the traffic received by the first security group 36a between the first, third and fourth firewall processing engines.

The directing the traffic from the first security group 36a step preferably includes the step of rebalancing traffic from the first security group 36a to the third and fourth firewall processing engines when the first firewall processing engine 24 fails. Preferably, after the step of transferring traffic to the first destination 20, there is the step of connecting a fifth firewall processing engine to the switch 14.

In the operation of the invention, the system 10 allows the network 12 owner to specify a subset of traffic which needs to be processed by a set of external processing devices, cull that traffic out, send it over to the devices, retrieve it, and forward it (now processed) to its destination. The culling can be done on any set of parameters inside the packet header.

This makes for a remarkably flexible and scalable packet processing system 10. This will be useful for banks, anyone who needs high-speed security, it can be used for massive encryption (1 Gbps and above), and a number of other applications that require the efforts of a number of CPU's to be grouped together. This system 10 basically acts as the taskmaster for a parallel processing platform.

A firewall processing engine 32, which is a combination of a hardware platform and a special software program, is a security computer that is normally implemented by situating the firewall processing engine 32 "in-line" between the internal and external networks. In these implementations, data traffic must pass through and be approved by the relatively slow hardware/software combination in order to pass through to the internal network. The system 10 moves the firewall processing engine 32 from the normal in-line position to a port of a special switching router that is in-line.

In figure 2, there are two security groups 36 configured on the switch 14 (A & B). Any traffic coming in those groups are pre-filtered (based on user-input filters), and sent to the firewall processing engines 32 for inspection. Traffic is load-shared across the engines based on a combination of input port and L2 parameters (L2 parameters are well known to one skilled in the art).

Once the traffic is redirected to the appropriate firewall processing engine, that particular engine will perform its programmed operations and returns the packet back to the switch 14, whereupon the switch 14 will forward it to its final destination.

If one of the firewall processing engines 32 fails, then the switch 14 will re-balance the traffic across the remaining engines. Currently, a failure is considered to be "link-down". More sophisticated failure detection algorithms are easy to implement with this system 10.

Figure 3 shows a basic configuration of the switching router 14, firewall processing engine 32, and external and internal networks A, B, respectively. As shown in figure 3, the firewall processing engine 32 is not in-line between the two networks (A & B).

The switching router 14 includes customized application specific integrated circuits ("ASICs"). These ASICs are designed to detect certain types of data packets that signify the opening and closing of a particular connection, as is well known in the art, and to then route them to the firewall processing engine 32. The firewall processing engine 32 verifies whether a particular connection request from external point A should be permitted, and if so, redirects the opening packet to the router for transmission to its intended destination on internal network B.

The system 10 is particularly well suited for use in connecting an internal network to external networks that communicate using the Transmission Control Protocol ("TCP"), which is the protocol used to carry data packets over the Internet. TCP is a connection-oriented protocol, meaning that before any communication can occur between two endpoints, a logical connection must be established so that both endpoints can expect the traffic. Data packets will not be accepted until the logical connection has been established. TCP uses certain opening and closing packets, known respectively as "SYN" and "FIN" packets. The first packet in a TCP communication session is the "SYN" packet and the last packet is the "FIN" packet. Each packet also includes the source and destination addresses of the two endpoints, as well as information that identifies the type of service that is being requested.

Referring now to figure 3, the firewall processing engine 32 operates as follows. A communication session begins when a computer on external network A sends a TCP open (SYN) packet to port 1 (P1) of the switching router 14. The switching router 14 includes the special ASICs that are capable of detecting the SYN and FIN packets. The router 14 detects the SYN packet from port 1 and redirects it to port 3 (P3), which is connected to the input of the software firewall processing engine 32. The firewall processing engine 32 examines the SYN packet, its associated source/destination addressing information, its requested service type, etc., in order to determine if the connection is permitted. The software firewall processing engine 32 would also include the ability to allow or not allow particular connections for many reasons, such as not permitting connections from particular domains, not permitting connections for particular services, etc., as is well known in the art.

If the connection is allowed, the SYN packet is routed back to the switching router 14 via port 4 (P4), and the router 14 then redirects the SYN packet to endpoint B in order to open the connection. If the connection request was not allowed by the software firewall processing engine 32, then the SYN packet would not have been redirected to B, and no connection would be established. After the computer at endpoint B receives the redirected SYN packet, it transmits an ACK (acknowledge) packet to endpoint A via the router, thereby indicating that a connection has been established. Normal packet traffic can then flow between points A and B, directly through the router 14, without being redirected to the software firewall processing engine 32.

The software firewall processing engine 32 maintains a list of the allowed and presently established connections. When the FIN packet for a particular connection is detected by the router 14, it is redirected to the software firewall processing engine 32, so that the connection can be properly cleared from the list of allowed connections. The FIN packet is then forwarded to endpoint B in order to close out the connection.

In the system 10, there is no communication channel established between the software firewall processing engine 32 and any external computer. The firewall processing engine 32 only communicates with the switching router by receiving SYN packets from port 3 (P3) for approval and transmitting those that are approved back to the router 14 via port 4 (P4). Because the system 10 is designed to be completely transparent to the external network, there would be no need to create such a connection between the firewall processing engine 32 and any external system.

The engines do not have to be firewall engines -- they can be any sort of traffic inspection/modification algorithm that the user can dream up. The system 10 specifies the ways and means of getting the traffic to and from that bank of processors.

The system's 10 load-sharing/redundancy scheme solves both the performance problem (by spreading the load across multiple CPUs) and the redundancy problem (by implementing a failover scheme).

There are two "physical" components and five "functional" components to the systems 10. The physical components are an exponeNT switching platform and a Check Point FireWall-1 firewall. The combination of these two physical components (with some software to hold them together) provides the following functional components: Packet Filtering, TCP Connection Inspection, UDP re-direction and software firewalling, and ICMP re-direction and software firewalling. The whole system 10 provides gigabit speed TCP/IP firewalling, line-rate switching and routing, strong QoS, ease of manageability, and strong resilience against failure.

Packet Filtering

There are two aspects to packet filtering; generic filters for user configuration, and security specific filters for filtering out and dealing with certain types of attacks.

Generic Packet Filters

Generic filters differ from filters on the traditional software-based routers in that implementation of these filters incurs no performance penalty. The following is an (edited) note from CERT, incorporated by reference herein, which outlines a reasonable filtering schema which will be used as the suggested base for the system 10:

> "The CERT staff encourages system managers, site network managers, and regional network providers to take the time to understand packet filtering issues. Because of the flaws in several TCP/IP services, a site must be able to restrict external access to these services. Sites should consider purchasing programmable routers. Network providers should offer packet filtering as a service option. Because of flaws in the protocol or chronic

system administration problems, we recommend that the following services be filtered:

- DNS zone transfers - socket 53(TCP)

  We suggest that sites filter socket 53(TCP) to prevent domain name service zone transfers. Permit access to socket 53(TCP) only from known secondary domain name servers. This prevents intruders from gaining additional knowledge about the systems connected to your local network

- tftpd - socket 69 (UDP)

  We have handled incidents that involved automated TFTP attempts. Many of the systems affected were using the TFTP daemon to boot other devices. Filtering TFTP connections would have protected the sites from this attack.

- link - socket 87 (TCP) (commonly used by intruders)

- SunRPC & NFS - socket 111 and 2049 (UDP and TCP)

- BSD UNIX "r" commands - sockets 512, 513, and 514 (TCP)

- lpd - socket 515 (TCP)

- uucpd - socket 540 (TCP)

- openwindows - socket 2000 (UDP and TCP)

- X windows - socket 6000+ (UDP and TCP)

  The X windows sockets range from socket 6000 to 6000 plus the highest number of X terminals on the same host. If your site does not need to provide other services to external users, those other services should be filtered. For example, filter telnet connections when all staff members are in the office, and filter FTP connections to all systems except to public information servers. In addition to filtering specific services, we recommend that sites also filter based on the source address field of the packets to prevent IP spoofing. More information on this technique can be found in CERT advisory CA-95:01, "IP Spoofing Attacks and Hijacked Terminal Connections," available by anonymous FTP

- To prevent denial of service attacks based on ICMP bombs, filter ICMP redirect and ICMP destination unreachable packets.

- In addition, sites should filter source-routed packets."

Blindly implementing this filtering in a LAN would cause all sorts of problems. For example, it is clearly not desirable to filter all Ipd traffic on a LAN where you're expecting to use UNIX printing services. However, a truly secure site might desire to restrict that traffic to those LAN segments that really are going to be using the service. Same with NFS, openwindows, and all of the other services mentioned. This is merely a generic overview of how filtering can help by closing certain security holes.

Security Specific Filtering

There are some attacks that switching hardware needs to deal with directly, either to performance-augment the software firewall, or because using the software firewall in the BNI configuration will disable the software firewall's ability to defend against these attacks. These are not meant to replace (nor will they ever be meant to replace) the software firewall, they are meant to augment its functionality. These filters are as follows:

- LAND attack filters (drop Source IP = Dest IP packets)
- Teardrop attack filters (drop Fragment offset = 1 packets)

- IP spoofing filters (2 types; 1 = check reverse path forwarding table; 2 = check manually input list of valid source IP addresses for a port)
- SMURF attack filters

TCP Connection Inspection

TCP Connection Inspection (also referred to as TCP cut-through) is probably the most interesting part of the LSS. Many TCP attacks against services which cannot simply be filtered out can be defended against by a software based firewall that keeps state (vs. a stateless filter). The most intuitive way of implementing this is to simply place a software firewall in line with every "protected" port. Unfortunately, there are still performance issues with this even with the fastest software firewalls on the market (e.g. Gigabit Ethernet connections are unprotectable).

The system 10 redirects the TCP connection packets to the software firewall, and allows the firewall to make the decision to forward the packet. Note that these are only the packets related to the actual opening of the connection; all other packets are forwarded as per their header information. If those packets aren't forwarded, then the connection never opens, and any subsequent packets would be dropped by the endstation. This functionality is exactly how the Check Point FireWall-1 product (available for purchase from Check Point), incorporated by reference herein, works when put in fastpath mode. The system 10 is taking the burden of forwarding the packets off of FireWall-1 and placing it in hardware.

This means that for long-lived, high volume connections, the performance can be much greater than with a software forwarding agent (as is the norm in firewalling.) The flip side of this is that the performance is tied to the "connection bandwidth" of the software firewall.

The Check Point FW-1 Software firewall is the firewall of choice for the system 10. Testing of FireWall-1 on a PII/400MHz machine indicates that FireWall-1 can support about 8000 connections/second when not burdened with forwarding the subsequent connection traffic.

Note that if there is only one firewall, and that firewall breaks, the customer will lose TCP connectivity on all of the firewalled ports 34. This is clearly unacceptable. There are also situations where the customer is going to need more than 8000 connections/second.

Fortunately, there is a mechanism by which to achieve both of these goals - redundancy and load balancing across multiple firewalls. The system 10 is implementing a trunking mechanism for use by inter-switch links. Any traffic coming in a certain port will be sent out a certain port in a trunk group. The trunk group, instead of terminating at another switch 14, can terminate on a group of firewalls, one firewall per trunk group link. There are multiple load-sharing mechanisms in place, but their purpose is identical - share load and implement redundancy.)

FireWall-1, on its side, can synchronize state between multiple firewalls. E.g. any connection that is cached in its connection table will be cached in the other firewalls within 50ms. While this turns out to not be relevant for TCP (when running in fastpath mode) (see below), it is very relevant for UDP. The concept of a group of firewalls all serving the same set of ports 34 is known as a "firewall pool".

FireWall-1 has a normal mode and fastpath mode. The normal mode inspects every packet, and caches connection state for every connection (or flow, TCP or otherwise). When in normal mode, all of the relevant TCP state packets (SYN, FIN, and ACK) need to be directed to the firewall. The FireWall-1 fastpath mode relies on the control mechanisms of TCP to correctly setup and tear down the connections, and so, only the SYN packets need to be directed to the firewall. (Which, when in fastpath mode, never caches connection state.) If it is desired to defend against SYN flooding attacks, the 1st ACK packet (non SYN/ACK) must also be directed to the firewall (this functionality is currently under investigation).

Since defending against SYN flooding requires ACK packets to be forwarded to the firewall system 10, it raises the question of sending the ACK packet to the firewall on which the original SYN was sent. This is not necessary, so the ACK packet can be sent to any firewall within the group of firewalls on the trunk links, and the FW-1 synchronization system 10 will deal with updating the state on the rest of the firewalls.

With FTP, it is desirable to have a mode whereby all traffic on the control channel is redirected to the software firewalls. The data channel is thence treated as a regular TCP connection. As the TCP control channel has its own dest TCP port, it is easy to segregate this traffic out and send it to the firewall system 10. The same holds true for SMTP traffic - all of it is currently redirected to FireWall-1 for full inspection (and content vectoring if relevant). This is the default behavior of the system 10.

The implementation of generic TCP/UDP filters on the exponeNT hardware allows for any application to be either cut-through without any inspection by FireWall-1 or fully inspected by FireWall-1. The implementation of generic IP source and/or destination filters on the exponeNT hardware allows for the same behavior for traffic coming from or going to any IP host or network 12.

UDP Re-Direction and Firewalling

UDP, being connectionless, doesn't have the manipulability of TCP for connection blocking purposes. However, certain protocols (NFS, for example), are very UDP, and can reveal sensitive information. Since UDP traffic is a small percentage of total traffic on the LAN, all UDP traffic can be sent through the firewalls. The use of firewall pools becomes very important for high-bandwidth UDP traffic.

FW-1 doesn't differentiate between fastpath and slowpath for UDP traffic. All traffic is inspected, and thus entries are made in the connection tables for each of the flows. Thus, synchronization between firewalls is important, and the number of flows that can be handled concurrently is also important. The number that was cited to me from Checkpoint is that FW-1 can handle approx. 25,000 concurrent connections.

ICMP Re-Direction and Firewalling

ICMP traffic would be handled in the same way as UDP traffic.

Configuration and Management

The system 10 requires configuration of both the exponeNT hardware (through the eVision management platform) and the FireWall-1 software.

One configuration that the system 10 uses is the "security group". Any set of ports 34 can be configured into a security group 36. A security group 36 shares a single firewall pool (and vice-versa, a single firewall pool only serves a single security group 36). Up to 16 security groups 36 can be configured on the switch 14, and up to 4 firewalls per firewall-pool. This allows for massive scalability and flexibility in performance and security. Traffic can pass between security groups 36 (if allowed by the policies).

One non-intuitive bit of configuration is that all security is applied at the input port level, e.g. traffic coming in port 1A1 is firewalled. This means that one must place all of the ports 34 that might have "bad-guys" attached to them into the relevant security groups 36.

Although the invention has been described in detail in the foregoing embodiments for the purpose of illustration, it is to be understood that such detail is solely for that purpose and that variations can be made therein by those skilled in the art without departing from the spirit and scope of the invention except as it may be described by the following claims.

WHAT IS CLAIMED IS:

1. A secure telecommunications system comprising:

a network on which traffic travels;

a switch connected to the network;

a first inspection engine connected to the switch, which receives traffic from the switch, processes the traffic to determine whether it is desired traffic or undesired traffic, which prevents undesired traffic from passing through it and which sends desired traffic back to the switch;

a second inspection engine connected to the switch, which receives traffic from the switch, processes the traffic to determine whether it is desired traffic or undesired traffic, which prevents undesired traffic from passing through it and which sends desired traffic back to the switch;

a first destination connected to the switch which receives desired traffic from the switch that has been processed by the first inspection engine; and

a second destination connected to the switch which receives desired traffic from the switch that has been processed by the second inspection engine.

2. A system as described in Claim 1 wherein the first inspection engine includes a first firewall processing engine and the second inspection engine includes a second firewall processing engine.

3. A system as described in Claim 2 wherein the switch has a first port and a second port connected to the network which receives traffic from the network, said switch directing traffic received at the first port to the first firewall processing engine and directing traffic received at the second port to the second firewall processing engine.

4. A system as described in Claim 3 including N additional firewall processing engines connected to the switch besides the first firewall processing engine and the second firewall processing engine so there are a total of N+2 firewall processing engines, where N is greater than or equal to 1 and is an integer.

5. A system as described in Claim 4 wherein the switch has N additional ports besides the first port and the second port, wherein each port is connected to a corresponding firewall processing engine.

6. A system as described in Claim 5 wherein the switch is configured into security groups with at least one of the N+2 firewall processing engines serving each security group.

7. A system as described in Claim 6 wherein the switch load-shares traffic for each security group across corresponding firewall processing engines serving the corresponding security group.

8. A system as described in Claim 7 wherein the switch rebalances traffic for a security group when one of the firewall processing engines serving the security group fails across the other firewall processing engines serving the security group.

9. A system as described in Claim 8 wherein the switch is scalable to allow for adding firewall processing engines.

10. A system as described in Claim 9 wherein the traffic includes bits and wherein the firewall processing engines serving a first security group of the security groups encrypt greater than 1 Gbps of traffic.

11. A system as described in Claim 10 wherein the network includes the Internet, and the first destination includes a first web server and the second destination includes a second web server.

12. A system as described in Claim 11 wherein the Internet includes a LAN.

13. A method for sending traffic over a secure telecommunications system comprising the steps of:

receiving traffic from a network at a switch connected to the network;

directing traffic to a first inspection engine connected to the switch and to a second inspection engine connected to the switch;

receiving traffic at the first inspection engine;

processing traffic received at the first inspection engine to determine whether it is desired traffic or undesired traffic;

sending the desired traffic back to the switch from the first inspection engine and discarding undesired traffic from the first inspection engine;

transferring desired traffic received by the switch from the first inspection engine to a first destination;

processing traffic received at the second inspection engine to determine whether it is desired traffic or undesired traffic;

sending the desired traffic back to the switch from the second inspection engine and discarding undesired traffic from the second inspection engine; and

transferring desired traffic received by the switch from the second inspection engine to a second destination.

14. A method as described in Claim 13 wherein the first and second inspection engines include a first firewall processing engine and a second firewall processing engine, respectively, and wherein the directing traffic step includes the step of directing traffic to the first firewall processing engine and second firewall processing engine and to a third firewall processing engine and a forth firewall processing engine.

15. A method as described in Claim 14 wherein the switch is configured into a first security group and a second security group, and the receiving step includes the step of receiving traffic at the first security group.

16. A method as described in Claim 15 wherein the directing step includes the step of directing the traffic from the first security group of the switch to the first, third and fourth firewall processing engines which serve the first security group of the switch, and directing traffic to the second firewall processing engine serving the second security group of the switch.

17. A method as described in Claim 16 wherein the receiving step includes the step of receiving traffic from the first security group at a first port of the switch and receiving traffic for the second security group at a second port of the switch.

18. A method as described in Claim 17 wherein the directing the traffic from the first security group includes the step of load-sharing by the switch the traffic received by the first security group between the first, third and fourth firewall processing engines.

19. A method as described in Claim 18 wherein the directing the traffic from the first security group includes the step of rebalancing traffic from the first security group to the third and fourth firewall processing engines when the first firewall processing engine fails.

20. A method as described in Claim 19 wherein after the step of transferring traffic to the first destination, there is the step of connecting a fifth firewall processing engine to the switch.

ABSTRACT OF THE DISCLOSURE

HARDWARE BASED SECURITY GROUPS, FIREWALL LOAD SHARING, AND FIREWALL REDUNDANCY

A secure telecommunications system. The system includes a network on which traffic travels. The system includes a switch connected to the network. The system includes a first inspection engine connected to the switch, which receives traffic from the switch, processes the traffic to determine whether it is desired traffic or undesired traffic, which prevents undesired traffic from passing through it and which sends desired traffic back to the switch. The system includes a second inspection engine connected to the switch, which receives traffic from the switch, processes the traffic to determine whether it is desired traffic or undesired traffic, which prevents undesired traffic from passing through it and which sends desired traffic back to the switch. The system includes a first destination connected to the switch which receives desired traffic from the switch that has been processed by the first inspection engine. The system includes a second destination connected to the switch which receives desired traffic from the switch that has been processed by the second inspection engine. A method for sending traffic over a secure telecommunications system.

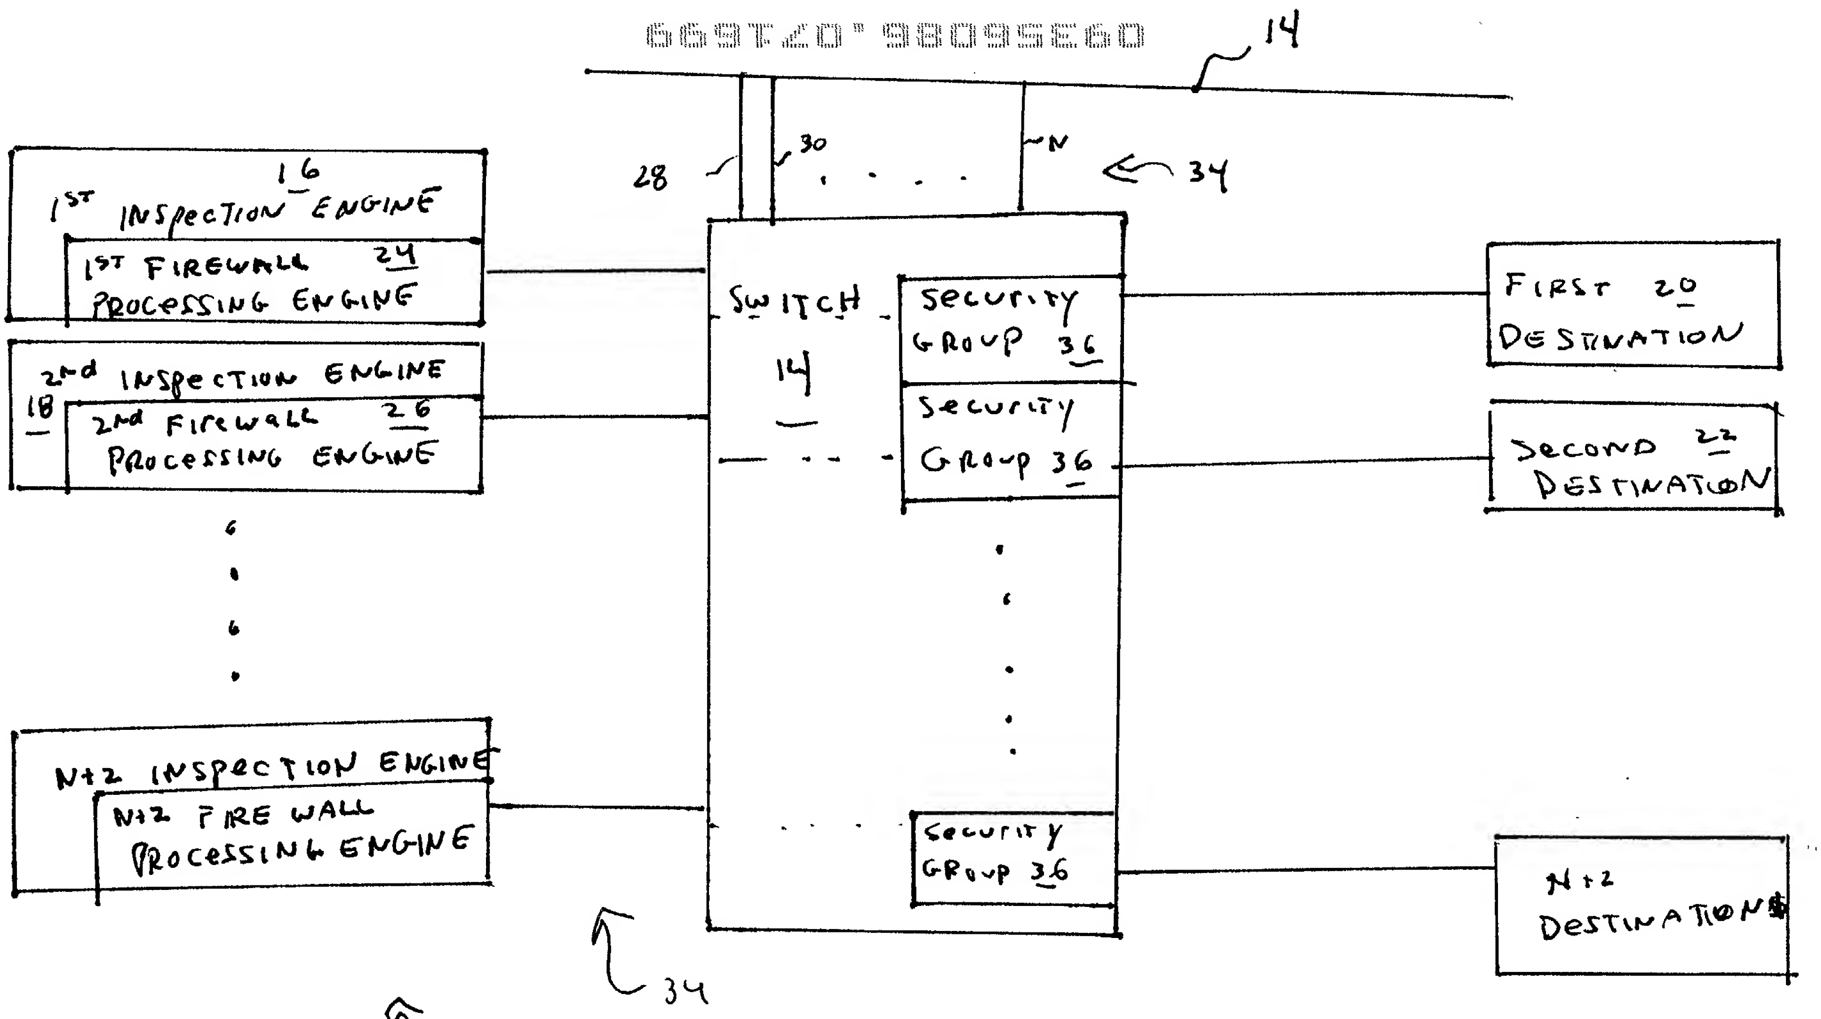
14
28
30
N
34
1ST INSPECTION ENGINE 16
1ST FIREWALL PROCESSING ENGINE 24
2nd INSPECTION ENGINE
18
2nd FIREWALL PROCESSING ENGINE 26
N+2 INSPECTION ENGINE
N+2 FIRE WALL PROCESSING ENGINE
SWITCH 14
SECURITY GROUP 36
SECURITY GROUP 36
SECURITY GROUP 36
FIRST 20 DESTINATION
SECOND 22 DESTINATION
N+2 DESTINATION
34
32

FIG 1

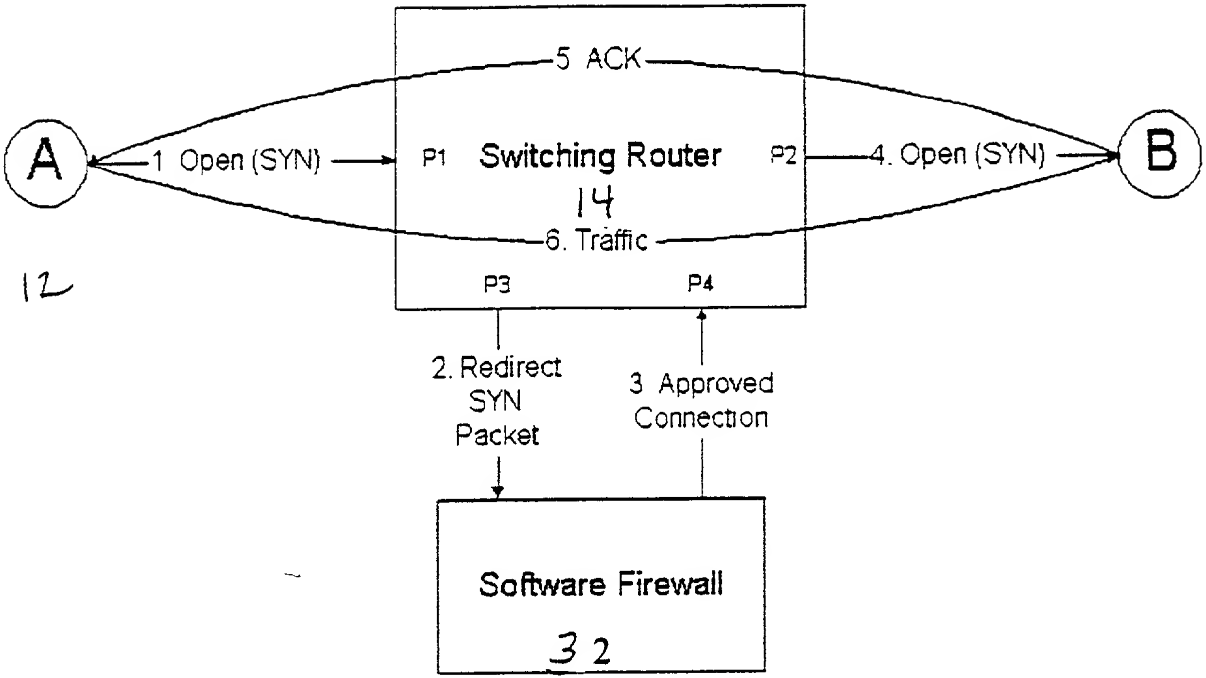
A
1. Open (SYN)
P1
Switching Router
14
P2
4. Open (SYN)
B
5 ACK
6. Traffic
12
P3
P4
2. Redirect SYN Packet
3 Approved Connection
Software Firewall
32

FIG 3

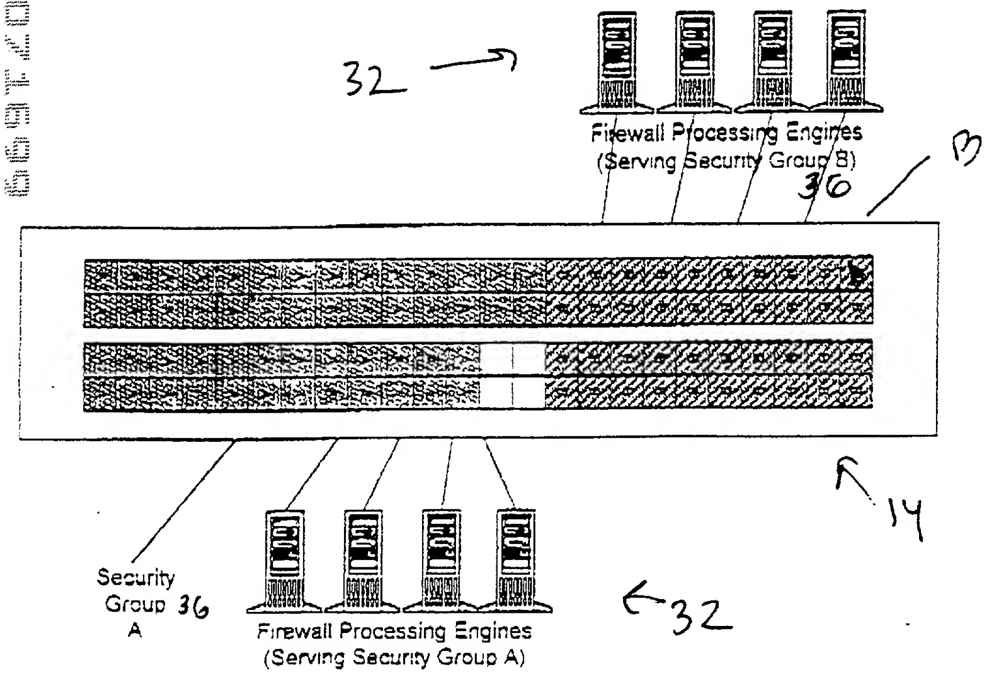
32
Firewall Processing Engines
(Serving Security Group B)
36
B
14
Security Group 36 A
Firewall Processing Engines
(Serving Security Group A)
32

FIG 2